arXiv:1411.0948v2 [cs.DC] 16 Sep 2015

On the push&pull protocol for rumour spreading*

Hüseyin Acan
School of Mathematical Sciences, Monash University
huseyin.acan@monash.edu

Andrea Collevecchio†
School of Mathematical Sciences, Monash University, and
Ca' Foscari University, Venice
andrea.collevecchio@monash.edu

Abbas Mehrabian‡
Department of Combinatorics and Optimization, University of Waterloo
amehrabi@uwaterloo.ca

Nick Wormald§
School of Mathematical Sciences, Monash University
nick.wormald@monash.edu

Abstract

The asynchronous push&pull protocol, a randomized distributed algorithm for spreading a rumour in a graph G, is defined as follows. Independent exponential clocks of rate 1 are associated with the vertices of G, one to each vertex. Initially, one vertex of G knows the rumour. Whenever the clock of a vertex x rings, it calls a random neighbour y: if x knows the rumour and y does not, then x tells y the rumour (a push operation), and if x does not know the rumour and y knows it, y tells x the rumour (a pull operation). The average spread time of G is the expected time it takes for all vertices to know the rumour, and the guaranteed spread time of G is the smallest time t such that with probability at least $1 - 1/n$, after time t all vertices know the rumour. The synchronous variant of this protocol, in which each clock rings precisely at times $1, 2, \ldots$, has been studied extensively.

We prove the following results for any n-vertex graph: In either version, the average spread time is at most linear even if only the pull operation is used, and the guaranteed spread time is within a logarithmic factor of the average spread time, so it is $O(n \log n)$. In the asynchronous version, both the average and guaranteed spread times are $\Omega(\log n)$. We give examples of graphs illustrating that these bounds are best possible up to constant factors.

We also prove the first analytical relationships between the guaranteed spread times in the two versions. Firstly, in all graphs the guaranteed spread time in the asynchronous version

*A preliminary version of this paper has appeared in proceedings of the ACM Symposium on Principles of Distributed Computing (PODC 2015), pages 405–412.

†Supported by ARC Discovery Project grant DP140100559 and ERC STREP project MATHEMACS.

‡Supported by the Vanier Canada Graduate Scholarships program.

§Supported by Australian Laureate Fellowships grant FL120100125.

is within an $O(\log n)$ factor of that in the synchronous version, and this is tight. Next, we find examples of graphs whose asynchronous spread times are logarithmic, but the synchronous versions are polynomially large. Finally, we show for any graph that the ratio of the synchronous spread time to the asynchronous spread time is $O(n^{2/3})$.

1 Introduction

Randomized rumour spreading is an important primitive for information dissemination in networks and has numerous applications in network science, ranging from spreading information in the WWW and Twitter to spreading viruses and diffusion of ideas in human communities. A well studied rumour spreading protocol is the *(synchronous) push&pull protocol*, introduced by Demers, Greene, Hauser, Irish, Larson, Shenker, Sturgis, Swinehart, and Terry [5] and popularized by Karp, Schindelhauer, Shenker, and Vöcking [23]. Suppose that one node in a network is aware of a piece of information, the 'rumour', and wants to spread it to all nodes quickly. The protocol proceeds in rounds. In each round, every *informed* node contacts a random neighbour and sends the rumour to it ('pushes' the rumour), and every *uninformed* nodes contacts a random neighbour and gets the rumour if the neighbour knows it ('pulls' the rumour).

A point to point communication network can be modelled as an undirected graph: the nodes represent the processors and the links represent communication channels between them. Studying rumour spreading has several applications to distributed computing in such networks, of which we mention just two. The first is in broadcasting algorithms: a single processor wants to broadcast a piece of information to all other processors in the network (see [20] for a survey). There are at least four advantages to the push&pull protocol: it puts much less load on the edges than naive flooding, it is simple (each node makes a simple local decision in each round; no knowledge of the global topology is needed; no state is maintained), scalable (the protocol is independent of the size of network: it does not grow more complex as the network grows) and robust (the protocol tolerates random node/link failures without the use of error recovery mechanisms, see [11]). A second application comes from the maintenance of databases replicated at many sites, e.g., yellow pages, name servers, or server directories. There are updates injected at various nodes, and these updates must propagate to all nodes in the network. In each round, a processor communicates with a random neighbour and they share any new information, so that eventually all copies of the database converge to the same contents. See [5] for details. Other than the aforementioned applications, rumour spreading protocols have successfully been applied in various contexts such as resource discovery [19], distributed averaging [4], data aggregation [24], and the spread of computer viruses [2].

In this paper we only consider simple, undirected and connected graphs. Given a graph and a starting vertex, the *spread time* of a certain protocol is the time it takes for the rumour to spread in the whole graph, i.e. the time difference between the moment the protocol is initiated and the moment when everyone learns the rumour. For the synchronous push&pull protocol, it turned out that the spread time is closely related to the *expansion profile* of the graph.

Let $\Phi(G)$ and $\alpha(G)$ denote the conductance and the vertex expansion of a graph G, respectively. After a series of results by various scholars, Giakkoupis [16, 17] showed the spread time is $O\left(\min\{\Phi(G)^{-1} \cdot \log n, \alpha(G)^{-1} \cdot \log^2 n\}\right)$. This protocol has recently been used to model news propagation in social networks. Doerr, Fouz, and Friedrich [6] proved an upper bound of $O(\log n)$ for the spread time on Barabási-Albert graphs, and Fountoulakis, Panagiotou, and Sauerwald [14] proved the same upper bound (up to constant factors) for the spread time on Chung-Lu random graphs.

All the above results assumed a synchronized model, i.e. all nodes take action simultaneously at discrete time steps. In many applications and certainly in real-world social networks, this assumption is not very plausible. Boyd, Ghosh, Prabhakar, Shah [4] proposed an asynchronous time model with a continuous time line. Each node has its own independent clock that rings at the times of a rate 1 Poisson process. (Since the times between rings is an exponential random variable, we shall call this an *exponential clock*.) The protocol now specifies for every node what to do when its own clock rings. The rumour spreading problem in the asynchronous time model has so far received less attention. Rumour spreading protocols in this model turn out to be closely related to Richardson's model for the spread of a disease [9] and to first-passage percolation [21] with edges having i.i.d. exponential weights. The main difference is that in rumour spreading protocols each vertex contacts one neighbour at a time. So, for instance in the 'push only' protocol, the net communication rate outwards from a vertex is fixed, and hence the rate that the vertex passes the rumour to any one given neighbour is inversely proportional to its degree (the push&pull protocol is a bit more complicated). Hence, the degrees of vertices play a crucial role not seen in Richardson's model or first-passage percolation. However, on regular graphs, the asynchronous push&pull protocol, Richardson's model, and first-passage percolation are essentially the same process, assuming appropriate parameters are chosen. In this sense, Fill and Pemantle [12] and Bollobás and Kohayakawa [3] showed that a.a.s. the spread time of the asynchronous push&pull protocol is $\Theta(\log n)$ on the hypercube graph. Janson [22] and Amini, Draief and Lelarge [1] showed the same results (up to constant factors) for the complete graph and for random regular graphs, respectively. These bounds match the same order of magnitude as in the synchronized case. Doerr, Fouz, and Friedrich [8] experimentally compared the spread time in the two time models. They state that 'Our experiments show that the asynchronous model is faster on all graph classes [considered here].' However, a general relationship between the spread times of the two variants has not been proved theoretically.

Fountoulakis, Panagiotou, and Sauerwald [14] studied the asynchronous push&pull protocol on Chung-Lu random graphs with exponent between 2 and 3. For these graphs, they showed that a.a.s. after some constant time, $n - o(n)$ nodes are informed. Doerr, Fouz, and Friedrich [7] showed that for the preferential attachment graph (the non-tree case), a.a.s. all but $o(n)$ vertices receive the rumour in time $O\left(\sqrt{\log n}\right)$, but to inform all vertices a.a.s., $\Theta(\log n)$ time is necessary and sufficient. Panagiotou and Speidel [27] studied this protocol on Erdős-Renyi random graphs and proved that if the average degree is $(1 + \Omega(1)) \log n$, a.a.s. the spread time is $(1 + o(1)) \log n$.

1.1 Our contribution

In this paper we answer a fundamental question about the asynchronous push&pull protocol: what are the minimum and maximum spread times on an n-vertex graph? Our proof techniques yield new results on the well studied synchronous version as well. We also compare the performances of the two protocols on the same graph, and prove the first theoretical relationships between their spread times.

We now formally define the protocols. In this paper G denotes the ground graph which is simple and connected, and n counts its vertices, and is assumed to be sufficiently large.

Definition (Asynchronous push&pull protocol)**.** Suppose that an independent exponential clock of rate 1 is associated with each vertex of G. Suppose that initially, some vertex v of G knows a piece of information, the so-called *rumour*. The rumour spreads in G as follows. Whenever the clock of a vertex x rings, this vertex performs an 'action': it calls a random neighbour y; if x knows the rumour and y does not, then x tells y the rumour (a *push* operation), and if x does not know the rumour and y knows it, y tells x the rumour (a *pull* operation). Note that if both x and y know the rumour or neither of them knows it, then this action is useless. Also, vertices have no memory, hence x may call the same neighbour several consecutive times. The *spread time* of G starting from v, written $\mathrm{ST}_{\mathsf{a}}(G, v)$, is the first time that all vertices of G know the rumour. Note that this is a continuous random variable, with two sources of randomness: the Poisson processes associated with the vertices, and random neighbour-selection of the vertices. The *guaranteed spread time* of G, written $\mathrm{gst}_{\mathsf{a}}(G)$, is the smallest deterministic number t such that for every $v \in V(G)$ we have $\mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(G, v) > t\right] \leq 1/n$. The *average spread time* of G, written $\mathrm{ast}_{\mathsf{a}}(G)$, is the smallest deterministic number t such that for every $v \in V(G)$ we have $\mathbb{E}\left[\mathrm{ST}_{\mathsf{a}}(G, v)\right] \leq t$.

Definition (Synchronous push&pull protocol)**.** Initially some vertex v of G knows the rumour, which spreads in G in a round-robin manner: in each round $1, 2, \ldots$, all vertices perform actions simultaneously. That is, each vertex x calls a random neighbour y; if x knows the rumour and y does not, then x tells y the rumour (a *push* operation), and if x does not know the rumour and y knows it, y tells x the rumour (a *pull* operation). Note that this is a synchronous protocol, e.g. a vertex that receives a rumour in a certain round cannot send it on in the same round. The *spread time* of G starting from v, $\mathrm{ST}_{\mathsf{s}}(G, v)$, is the first time that all vertices of G know the rumour. Note that this is a discrete random variable, with one source of randomness: the random neighbour-selection of the vertices. The *guaranteed spread time* of G, written $\mathrm{gst}_{\mathsf{s}}(G)$, and the *average spread time* of G, written $\mathrm{ast}_{\mathsf{s}}(G)$, are defined in an analogous way to the asynchronous case.

We remark that the notion of 'guaranteed spread time' was first defined by Feige, Peleg, Raghavan and Upfal [11] under the name 'almost sure rumor coverage time' for the 'push only' protocol. (In this protocol, which was studied prior to push&pull, the informed nodes push the rumour, but the uninformed ones do nothing. The 'pull only' protocol is defined conversely.)

It turns out that changing the starting vertex affects the spread time by at most a multiplicative factor of 2. Specifically, in Proposition 14 we prove that for any two vertices u and

v, $\mathrm{ST}_{\mathsf{s}}(G,u) \overset{s}{\leq} 2\,\mathrm{ST}_{\mathsf{s}}(G,v)$ and $\mathrm{ST}_{\mathsf{a}}(G,u) \overset{s}{\leq} 2\,\mathrm{ST}_{\mathsf{a}}(G,v)$. (For random variables X and Y, $X \overset{s}{\leq} Y$ means X is *stochastically dominated by* Y, that is, for any t, $\mathbb{P}[X \geq t] \leq \mathbb{P}[Y \geq t]$.) These imply $\mathrm{ast}_{\mathsf{s}}(G) \leq 2\mathbb{E}[\mathrm{ST}_{\mathsf{s}}(G,v)]$ and $\mathrm{ast}_{\mathsf{a}}(G) \leq 2\mathbb{E}[\mathrm{ST}_{\mathsf{a}}(G,v)]$ for any vertex v.

Our first main result is the following theorem.

Theorem 1. *For large enough n, the following hold for any n-vertex graph G.*

$$(1-1/n)\,\mathrm{ast}_{\mathsf{a}}(G) \leq \mathrm{gst}_{\mathsf{a}}(G) \leq e\,\mathrm{ast}_{\mathsf{a}}(G)\log n\,, \tag{1}$$

$$\frac{1}{5}\log n < \mathrm{ast}_{\mathsf{a}}(G) < 4n\,, \tag{2}$$

$$\frac{1}{5}\log n \leq \mathrm{gst}_{\mathsf{a}}(G) \leq 4en\log n\,. \tag{3}$$

Moreover, these bounds are asymptotically best possible, up to the constant factors.

Our proof of the right-hand bound in (2) is based on the pull operation only, so this bound applies equally well to the 'pull only' protocol.

The arguments for (1) and the right-hand bounds in (2) and (3) can easily be extended to the synchronous variant, giving the following theorem. The bound (6) below also follows from [10, Theorem 4], but here we also show its tightness.

Theorem 2. *The following hold for any n-vertex graph G.*

$$(1-1/n)\,\mathrm{ast}_{\mathsf{s}}(G) \leq \ \mathrm{gst}_{\mathsf{s}}(G)\ \leq e\,\mathrm{ast}_{\mathsf{s}}(G)\log n\,, \tag{4}$$

$$\mathrm{ast}_{\mathsf{s}}(G) < 4.6n\,, \tag{5}$$

$$\mathrm{gst}_{\mathsf{s}}(G) < 4.6en\log n\,. \tag{6}$$

Moreover, these bounds are asymptotically best possible, up to the constant factors.

Open problem 3. Find the best possible constants factors in Theorems 1 and 2.

We next turn to studying the relationship between the asynchronous and synchronous variants on the same graph. Let $H_n := \sum_{i=1}^{n} 1/i$ denote the nth harmonic number. It is well known that $H_n = \log n + O(1)$.

Theorem 4. *For any G we have $\mathrm{ast}_{\mathsf{a}}(G) \leq H_n \times \mathrm{ast}_{\mathsf{s}}(G)$ and $\mathrm{gst}_{\mathsf{a}}(G) \leq 8\,\mathrm{gst}_{\mathsf{s}}(G)\log n$, and these bounds are best possible, up to the constant factors.*

For all graphs we examined stronger results hold, which suggests the following conjecture.

Conjecture 5. For any n-vertex graph G we have $\mathrm{ast}_{\mathsf{a}}(G) \leq \mathrm{ast}_{\mathsf{s}}(G) + O(\log n)$ and $\mathrm{gst}_{\mathsf{a}}(G) \leq \mathrm{gst}_{\mathsf{s}}(G) + O(\log n)$.

Our last main result is the following theorem, whose proof is somewhat technical, and uses couplings with the sequential rumour spreading protocol.

Theorem 6. *For any $\alpha \in [0,1)$ we have*

$$\mathrm{gst}_{\mathsf{s}}(G) \leq n^{1-\alpha} + 64\mathrm{gst}_{\mathsf{a}}(G)n^{(1+\alpha)/2}\,.$$

Corollary 7. *We have*

$$\frac{\mathrm{gst}_{\mathsf{s}}(G)}{\mathrm{gst}_{\mathsf{a}}(G)} = \Omega(1/\log n) \quad \text{and} \quad \frac{\mathrm{gst}_{\mathsf{s}}(G)}{\mathrm{gst}_{\mathsf{a}}(G)} = O\big(n^{2/3}\big)\,,$$

and the left-hand bound is asymptotically best possible, up to the constant factor. Moreover, there exist infinitely many graphs for which this ratio is $\Omega\left(n^{1/3}(\log n)^{-4/3}\right)$.

Open problem 8. What is the maximum possible value of the ratio $\mathrm{gst}_{\mathsf{s}}(G)/\mathrm{gst}_{\mathsf{a}}(G)$ for an n-vertex graph G?

We make the following conjecture.

Conjecture 9. For any n-vertex graph G we have

$$\frac{\mathrm{gst}_{\mathsf{s}}(G)}{\mathrm{gst}_{\mathsf{a}}(G)} = O\left(\sqrt{n}\,(\log n)^{O(1)}\right)\,,$$

and this is tight for infinitely many graphs.

The parameters $\mathrm{ast}_{\mathsf{s}}(G)$ and $\mathrm{ast}_{\mathsf{a}}(G)$ can be approximated easily using the Monte Carlo method: simulate the protocols several times, measuring the spread time of each simulation, and output the average. Another open problem is to design a *deterministic* approximation algorithm for any one of $\mathrm{ast}_{\mathsf{a}}(G)$, $\mathrm{gst}_{\mathsf{a}}(G)$, $\mathrm{ast}_{\mathsf{s}}(G)$ or $\mathrm{gst}_{\mathsf{s}}(G)$.

For the proofs we use standard graph theoretic arguments and well known properties of the exponential distribution and Poisson processes, in particular the memorylessness, and the fact that the union of two Poisson processes is another Poisson process. For proving Theorem 6 we define a careful coupling between the synchronous and asynchronous protocols.

Previous work on the asynchronous push&pull protocol has focused on special graphs. This paper is the first systematic study of this protocol on all graphs. We believe this protocol is fascinating and is quite different from its synchronous variant, in the sense that different techniques are required for analyzing it, and the spread times of the two versions can be quite different. Our work makes significant progress on better understanding of this protocol, and will hopefully inspire further research on this problem.

A collection of known and new bounds for the average spread times of many graph classes is given in Table 1. In Section 2 we prove some preliminary results and study some examples, which demonstrate tightness of some of the above bounds. Theorems 1 and 2 are proved in Section 3. Theorems 4 and 6 and Corollary 7 are proved in Section 4.

Graph G	$\mathrm{ast}_{\mathtt{s}}(G)$	$\mathrm{ast}_{\mathtt{a}}(G)$
Path	$(4/3)n + O(1)$	$n + O(1)$
Star	2	$\log n + O(1)$
Complete	$\sim \log_3 n$ [23]	$\log n + o(1)$ [22]
String of diamonds $\mathcal{S}_{m,k}$ (see Section 2.5)	$\boldsymbol{\Omega(m)}$	$\boldsymbol{O(\log n + m/\sqrt{k})}$
Hypercube	$\Theta(\log n)$ [11]	$\Theta(\log n)$ [12]
Random graphs $\mathcal{G}(n,p)$ $1 < \frac{np}{\log n}$ fixed	$\Theta(\log n)$ [11]	$\sim \log n$ [27]
Random d-regular graphs $2 < d$ fixed	$\Theta(\log n)$ [13]	$\sim (\log n)(d-1)/(d-2)$ [1]
Preferential attachment graphs (Barabási-Albert model)	$\Theta(\log n)$ [6]	$\Theta(\log n)$ [7]
Random geometric graphs in $[0, \sqrt[d]{n}]^d$ with edge threshold r above giant component threshold	$\Theta(\sqrt[d]{n}/r + \log n)$ [15]	$\boldsymbol{O\left(\log n \cdot \sqrt[d]{n}/r + \log^2 n\right)}$
Random k-trees ($2 \le k$ fixed)	$\Omega\left(n^{1/(k+3)}\right)$ [26]	$\Omega\left(n^{1/(k+3)}\right)$ [26]
General	$O\left(\Delta(G)(\mathrm{diam}(G) + \log n)\right)$ [11]	$O\left(\Delta(G)(\mathrm{diam}(G) + \log n)\right)$ [11]
General	$O\left((\log n)/\Phi(G)\right)$ [16]	$\boldsymbol{O\left((\log^2 n)/\Phi(G)\right)}$
General	$O\left((\log \Delta(G) \cdot \log n)/\alpha(G)\right)$ [17]	$\boldsymbol{O\left((\log \Delta(G) \cdot \log^2 n)/\alpha(G)\right)}$

Table 1: Average spread times of some graph classes are shown. For many of the entries, the relevant paper indeed proves an asymptotically almost sure bound for the spread time. Bold entries are new to this paper. Results for random graph classes hold *asymptotically almost surely* as the number of vertices grows. The notation $\sim$ means equality up to a $1 + o(1)$ factor. $\Delta(G)$ denotes the maximum degree of G. For $S \subseteq V(G)$, let ∂S be the set of vertices in $V(G) \setminus S$ that have a neighbour in S, and let $e(S, V(G) \setminus S)$ be the number of edges between S and $V(G) \setminus S$, and let $\mathtt{vol}(S) = \sum_{u \in S} \deg(u)$. Then $\alpha(G) := \min\left\{\frac{|\partial S|}{|S|} : S \subseteq V(G), 0 < |S| \le |V(G)|/2\right\}$ and $\Phi(G) := \min\left\{\frac{e(S,V(G)\setminus S)}{\mathtt{vol}(S)} : S \subseteq V(G), 0 < \mathtt{vol}(S) \le \mathtt{vol}(V(G))/2\right\}$.

2 Preliminaries and examples

Let us denote $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ and $\mathbb{N} = \{1, 2, \dots\}$. Let $\mathrm{Geo}(p)$ denote a geometric random variable with parameter p taking values in $\mathbb{N}_0$; namely for every $k \in \mathbb{N}_0$, $\mathbb{P}[\mathrm{Geo}(p) = k] = (1-p)^k p$. Let $\mathrm{Exp}(\lambda)$ denote an exponential random variable with parameter λ and mean $1/\lambda$. For random variables X and Y, $X \stackrel{d}{=} Y$ means X and Y have the same distribution. All logarithms are natural. For functions f and g, $f \sim g$ means $\lim f(n)/g(n) = 1$ as n grows.

We start by making a few observations valid of all graphs.

Observation 10. Consider the asynchronous variant. Let uv be an edge. Whenever v's clock rings, it calls u with probability $1/\deg(v)$. Hence, for each vertex v, we can replace v's clock by one exponential clock for each incident edge, these clocks being independent of all other clocks and having rate $1/\deg(v)$.

Observation 11. Whenever a new vertex is informed, by memorylessness of the exponential random variable, we may imagine that all clocks are restarted.

The following definition will be used throughout.

Definition (Communication time)**.** For an edge $e = uv$, the *communication time via edge* e, written $T(e)$, is defined as follows. Suppose τ is the first time that one of u and v learns the rumour, and ρ is the first time after τ that one of u and v calls the other one. Then $T(e) = \rho - \tau$, which is nonnegative. Note that after time ρ, both u and v know the rumour.

Observation 12. Let $uv \in E(G)$. In the synchronous version,

$$T(uv) \stackrel{d}{=} 1 + \min\{\mathrm{Geo}(1/\deg(u)), \mathrm{Geo}(1/\deg(v))\}\,.$$

Using Observations 10 and 11, we obtain a nicer formula for the asynchronous version.

Proposition 13. *Let $uv \in E(G)$. In the asynchronous version,*

$$T(uv) \stackrel{d}{=} \mathrm{Exp}(1/\deg(u) + 1/\deg(v))\,. \tag{7}$$

Moreover, the random variables $\{T_e\}_{e \in E(G)}$ are mutually independent.

Proof. By Observations 10 and 11, the $T(e)$'s are mutually independent, and moreover, $T(uv)$ is the minimum of two independent exponential random variables with rates $1/\deg(v)$ and $1/\deg(u)$. □

We next prove that changing the starting vertex affects the spread time by at most a multiplicative factor of 2.

Proposition 14. *For any two vertices u and v of G we have $\mathrm{ST}_{\mathsf{s}}(G,u) \stackrel{s}{\leq} 2\,\mathrm{ST}_{\mathsf{s}}(G,v)$ and also $\mathrm{ST}_{\mathsf{a}}(G,u) \stackrel{s}{\leq} 2\,\mathrm{ST}_{\mathsf{a}}(G,v)$.*

Proof. We first consider the synchronous protocol. Let $C(u,v)$ denote the first time that v learns the rumour, assuming initially only u knows it. We claim that

$$C(u,v) \stackrel{d}{=} C(v,u)\,, \tag{8}$$

which would imply

$$\mathrm{ST}_{\mathsf{s}}(G,u) \stackrel{s}{\leq} C(u,v) + \mathrm{ST}_{\mathsf{s}}(G,v) \stackrel{d}{=} C(v,u) + \mathrm{ST}_{\mathsf{s}}(G,v) \stackrel{s}{\leq} 2\,\mathrm{ST}_{\mathsf{s}}(G,v)\,.$$

In every round of an execution of the protocol, each vertex contacts a neighbour. We call this an *action*, and the *signature* of this action, is a function $a : V \to V$ mapping each vertex to a neighbour. Hence, m rounds of the protocol can be encoded as $(u, a_1a_2\cdots a_m)$, where u is the vertex knowing the rumour initially, and $a_1a_2\cdots a_m$ is a sequence of signatures. Let $I(u, a_1a_2\cdots a_m)$ denote the set of informed vertices after m rounds. Note that in each round, the signature of the action taken is a uniformly random one. Hence $\mathbb{P}[C(u,v) \leq k]$ equals the proportion of the signature-sequences $a_1a_2\cdots a_k$ of length k that satisfy $v \in I(u, a_1\cdots a_k)$. If $v \in I(u, a_1\cdots a_k)$, then looking at the (u,v)-path through which v was informed, we see that $u \in I(v, a_ka_{k-1}\cdots a_2a_1)$. Therefore, $\mathbb{P}[C(u,v) \leq k] = \mathbb{P}[C(v,u) \leq k]$ for any k, and this proves (8).

We now consider the asynchronous protocol. Let $D(u,v)$ denote the first time that v learns the rumour, assuming initially only u knows it. Again, it suffices to prove

$$D(u,v) \stackrel{d}{=} D(v,u)\,. \tag{9}$$

By Proposition 13, for any edge uv we have $T(uv) \stackrel{d}{=} \mathrm{Exp}(1/\deg(u) + 1/\deg(v))$. Moreover, the variables $\{T(e)\}_{e\in E}$ are mutually independent. We define a collection of mutually independent random variables $\{R(e)\}_{e\in E}$, such that for any edge uv,

$$R(uv) \stackrel{d}{=} \mathrm{Exp}(1/\deg(u) + 1/\deg(v))\,.$$

Let $\mathcal{P}$ denote the set of all (u,v)-paths. Then we have

$$D(u,v) = \min\left\{\sum_{e\in P} T(e) : P \in \mathcal{P}\right\} \stackrel{d}{=} \min\left\{\sum_{e\in P} R(e) : P \in \mathcal{P}\right\}\,.$$

By symmetry, $D(v,u)$ has exactly the same distribution, and (9) follows. □

We next study some important graphs and bound their spread times, partly for showing tightness of some of the bounds obtained, and partly to serve as an introduction to the behaviour of the protocols.

2.1 The complete graph

For the complete graph, K_n, by symmetry what matters at any time is not the actual set of informed vertices, but only the number of vertices that have the rumour. In the asynchronous case, by Proposition 13 and Observation 10, we can imagine a exponential clock for each edge, having

rate $2/(n-1)$ and independent of all other clocks. Let $T_1 = 0$ and denote by T_k the first time that there are k informed vertices. We can at this time simply restart all $k(n-k)$ clocks at edges joining informed to uninformed vertices (see Observation 11). When the next alarm rings, a new vertex receives the rumour. Thus, $T_{k+1} - T_k$ is distributed as the mimimum of $k(n-k)$ independent exponential random variables each with rate $2/(n-1)$, i.e. as $\mathrm{Exp}(2k(n-k)/(n-1))$. Hence by linearity of expectation,

$$\mathrm{ast}_{\mathsf{a}}(K_n) = \mathbb{E}\left[\mathrm{ST}_{\mathsf{a}}(K_n, v)\right] = \mathbb{E}\left[T_n\right] = \mathbb{E}\left[T_1\right] + \sum_{k=1}^{n-1} \mathbb{E}\left[T_{k+1} - T_k\right] = \sum_{k=1}^{n-1} \frac{n-1}{2k(n-k)}\,.$$

We have

$$\sum_{k=1}^{n-1} \frac{n-1}{2k(n-k)} = \left(\frac{n-1}{2n}\right) \sum_{k=1}^{n-1} \left\{\frac{1}{k} + \frac{1}{n-k}\right\} \sim \left(\frac{n-1}{2n}\right)(2\log n) \sim \log n\,,$$

so $\mathrm{ast}_{\mathsf{a}}(K_n) \sim \log n$. In fact, Janson [22, Theorem 1.1(ii)] showed that a.a.s. $\mathrm{ST}_{\mathsf{a}}(K_n, v) \sim \log n$. Moreover, by slightly altering his proof we get $\mathrm{gst}_{\mathsf{a}}(K_n) \sim (3/2)\log n$.

For the synchronous version, Karp et al. [23, Theorem 2.1] showed that $\mathrm{ST}_{\mathsf{s}}(K_n, v) \sim \log_3 n$ a.a.s. It follows that $\mathrm{ast}_{\mathsf{s}}(K_n) \sim \log_3 n$. It is implicit in their proof that $\mathrm{gst}_{\mathsf{s}}(K_n) = O(\log n)$.

2.2 The star

The star G_n^* with n vertices has $n-1$ leaves and a central vertex that is adjacent to every other vertex. It is clear that $\mathrm{ST}_{\mathsf{s}}(G_n^*, v) = 1$ if v is the central vertex and $\mathrm{ST}_{\mathsf{s}}(G_n^*, v) = 2$ otherwise. So we have $\mathrm{ast}_{\mathsf{s}}(G_n^*) = \mathrm{gst}_{\mathsf{s}}(G_n^*) = 2$. Below we will show that $\mathrm{ast}_{\mathsf{a}}(G_n^*) \sim \log n$ and $\mathrm{gst}_{\mathsf{a}}(G_n^*) \sim 2\log n$. This graph gives that the left-hand bounds in (1), (2), (3), (4) and Corollary 7, and Theorem 4, are tight, up to constant factors.

We now show that $\mathrm{ast}_{\mathsf{a}}(G_n^*) \sim \log n$. The intuition is that in the asynchronous case, the spread time is close to the time the last vertex makes its first call. By Proposition 13, all communication times are independent and distributed as $\mathrm{Exp}(n/(n-1))$. Let $X_1, \ldots, X_{n-1}$ be independent $\mathrm{Exp}(n/(n-1))$ random variables. Then

$$\mathrm{ST}_{\mathsf{a}}(G_n^*, v) \stackrel{d}{=} \begin{cases} \max\{X_1, \ldots, X_{n-1}\} & \text{if } v \text{ is the central vertex} \\ X_1 + \max\{X_2, \ldots, X_{n-1}\} & \text{if } v \text{ is a leaf.} \end{cases}$$

It follows that $\mathrm{ast}_{\mathsf{a}}(G_n^*) \sim \log n$ and a.a.s. $\mathrm{ST}_{\mathsf{a}}(G_n^*, v) \sim \log n$ for any v.

We finally show that $\mathrm{gst}_{\mathsf{a}}(G_n^*) \sim 2\log n$. Let $\lambda = n/(n-1)$, $Y = X_1$ and define $Z = \max\{X_2, \ldots, X_{n-1}\}$. Fix $\varepsilon \in (0,1)$. We have

$$\mathbb{P}\left[Z \le \log n\right] = (1 - e^{-\lambda \log n})^{n-2} \sim 1/e.$$

Hence

$$\mathbb{P}\left[Y + Z \ge (2-\varepsilon)\log n\right] \ge \mathbb{P}\left[Y \ge (1-\varepsilon)\log n\right] \mathbb{P}\left[Z \ge \log n\right] = \Theta(n^{-1+\varepsilon}). \tag{10}$$

Now let $a = (1+\varepsilon/2)\log n$ and let A be the event $\{Y \le a\}$. We have $\mathbb{P}[A^c] = O(n^{-1-\varepsilon/2})$. Also

$$\begin{aligned}\mathbb{P}[Y+Z \ge (2+\varepsilon)\log n] &\le \mathbb{P}[A^c] + \int_0^a \lambda e^{-\lambda y}\,\mathbb{P}[Z > (2+\varepsilon)\log n - y]\,\mathrm{d}y \\ &= O(n^{-(1+\varepsilon/2)}) + \int_0^a \lambda e^{-\lambda y}\,\mathbb{P}[Z > (2+\varepsilon)\log n - y]\,\mathrm{d}y. \end{aligned} \tag{11}$$

Using the independence of the X_i,

$$\begin{aligned}\mathbb{P}[Z > (2+\varepsilon)\log n - y] &= 1 - \left(1 - e^{-\lambda[(2+\varepsilon)\log n - y]}\right)^{n-2} \\ &\le 1 - \left(1 - e^{-\lambda[(2+\varepsilon)\log n - y]}\right)^{n} \\ &= 1 - \left(1 - \frac{e^{\lambda y}}{n^{\lambda(2+\varepsilon)}}\right)^{n} \le \frac{e^{\lambda y}}{n^{1+\varepsilon}}. \end{aligned} \tag{12}$$

The last inequality can be justified by expanding the left hand side and using the fact that $\lambda > 1$. Using (12) in (11) we get

$$\mathbb{P}[Y+Z \ge (2+\varepsilon)\log n] \le O(n^{-(1+\varepsilon/2)}) + \frac{1}{n^{1+\varepsilon}}\int_0^a \lambda e^{-\lambda y}\, e^{\lambda y} dy = O\left(n^{-(1+\varepsilon/2)}\right).$$

This equation together with equation (10) implies $\mathrm{gst}_{\mathsf{a}}(G_n^*) \sim 2\log n$.

2.3 The path

For the path graph P_n, we have $\mathrm{ast}_{\mathsf{a}}(P_n) \sim n$, which shows that the right-hand bound in (2) is tight, up to the constant factor. Moreover, $\mathrm{gst}_{\mathsf{a}}(P_n) \sim n$. For the synchronous protocol, we have $\mathrm{ast}_{\mathsf{s}}(P_n) = (4/3)n - 2$, which shows that the right-hand bound in (5) is tight, up to the constant factor. Finally, we have $\mathrm{gst}_{\mathsf{s}}(P_n) \sim (4/3)n$. Detailed calculations follow.

Label the vertices in the path as $(v_1, \ldots, v_n)$. In this graph, the spread times in the synchronous and asynchronous variants are close to each other. We first consider the asynchronous variant. Let e be an edge. By Proposition 13, if e connects two internal vertices, then $T(e) \stackrel{d}{=} \mathrm{Exp}(1)$, and otherwise, $T(e) \stackrel{d}{=} \mathrm{Exp}(3/2)$. Thus if the rumour starts from one of the endpoints, say v_1, we have

$$\mathrm{ST}_{\mathsf{a}}(P_n, v_1) \stackrel{d}{=} \sum_{i=1}^{n-1} X_i, \tag{13}$$

where X_i's are independent exponential random variables, X_1 and X_{n-1} with rates $3/2$ and the rest with rates 1. It follows that $\mathbb{E}[\mathrm{ST}_{\mathsf{a}}(P_n, v_1)] = (n-3) + 2(2/3) = n - 5/3$. With similar computations, it is easy to see that this is the worst case, i.e. $\mathrm{ast}_{\mathsf{a}}(P_n) = n - 5/3$.

Next we show $\mathrm{gst}_{\mathsf{a}}(P_n) \sim n$. Fix $\varepsilon > 0$. Note that $\sum_{i=2}^{n-2} X_i$ is a sum of i.i.d. random variables, hence by Cramér's Theorem (see, e.g., [18, Theorem 5.11.4]), the probability that it deviates by at least εn from its expected value is $\exp(-\Omega(n))$. Moreover, (13) means that $\mathrm{ST}_{\mathsf{a}}(P_n, v_1)$ is $\sum_{i=2}^{n-2} X_i$ plus two exponential random variables with constant rate, and the same statement is true for it as well, so $\mathrm{gst}_{\mathsf{a}}(P_n) \sim n$.

Now consider the synchronous case. Let e be an edge. By Observation 12, if e connects two internal vertices, then $T(e) \stackrel{d}{=} 1 + \mathrm{Geo}(3/4)$, and otherwise, $T(e) = 1$. Thus if the rumour starts from one of the endpoints, say v_1, we have

$$\mathrm{ST}_{\mathsf{s}}(P_n, v_1) \stackrel{d}{=} n - 1 + \sum_{i=2}^{n-2} X_i, \tag{14}$$

where X_i's are independent $\mathrm{Geo}(3/4)$ random variables. It follows that $\mathbb{E}\left[\mathrm{ST}_{\mathsf{s}}(P_n, v_1)\right] = (4/3)n - 2$. With similar computations, it is easy to see that this is the worst case, i.e. $\mathrm{ast}_{\mathsf{s}}(P_n) = (4/3)n - 2$. An argument similar to the one for the asynchronous variant gives $\mathrm{gst}_{\mathsf{s}}(P_n) \sim (4/3)n$.

2.4 The double star

Consider the tree DS_n consisting of two adjacent vertices of degree $n/2$ and $n-2$ leaves, see Figure 1(Left). Below we will show that $\mathrm{gst}_{\mathsf{a}}(DS_n)$ and $\mathrm{gst}_{\mathsf{s}}(DS_n)$ are both $\Theta(n \log n)$, while the average times $\mathrm{ast}_{\mathsf{a}}(DS_n)$ and $\mathrm{ast}_{\mathsf{s}}(DS_n)$ are $\Theta(n)$. This example hence shows tightness of the right-hand bounds in (1), (3), (4) and (6) up to constant factors. The main delay in spreading the rumour in this graph comes from the edge joining the two centres. The idea is that it takes $\Theta(n)$ units of time on average for this edge to pass the rumour, but to be sure that this has happened with probability $1 - 1/n$, we need to wait $O(n \log n)$ units of time. Detailed calculations follow.

Indeed we will show that $\mathrm{ast}_{\mathsf{a}}(DS_n)$ and $\mathrm{ast}_{\mathsf{s}}(DS_n)$ are asymptotic to $n/4$, and $\mathrm{gst}_{\mathsf{a}}(DS_n)$ and $\mathrm{gst}_{\mathsf{s}}(DS_n)$ are asymptotic to $n \log n/4$. First, consider the asynchronous case. Here, by Proposition 13, $T(e^*) = \mathrm{Exp}(4/n)$. So, the rumour passes from one centre to the other one in $n/4$ time units on average. On the other hand, the leaves learn the rumour in $\Theta(\log n)$ time on average, as in the star graph. Combining the two, we get $\mathrm{ast}_{\mathsf{a}}(DS_n) \sim n/4$.

For the guaranteed spread time, note that if $c < 1/4$ then

$$\mathbb{P}\left[T(e^*) \geq cn \log n\right] = \exp(-n/4 \times cn \log n) \geq 1/n\,.$$

Thus $\mathrm{gst}_{\mathsf{a}}(DS_n) \geq n \log n/4$. Straightforward calculations give that if $c > 1/4$ then for any vertex v, $\mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(DS_n, v) > cn \log n\right] < 1/n$, whence $\mathrm{gst}_{\mathsf{a}}(DS_n) \sim (n \log n)/4$.

In the synchronous case, for any v we have $T(e^*) + 1 \leq \mathrm{ST}_{\mathsf{s}}(DS_n, v) \leq T(e^*) + 2$ and by Observation 12,

$$T(e^*) \stackrel{d}{=} 1 + \min\{\mathrm{Geo}(2/n), \mathrm{Geo}(2/n)\} \stackrel{d}{=} 1 + \mathrm{Geo}(4/n - 4/n^2)\,,$$

hence $\mathrm{ast}_{\mathsf{s}}(DS_n) = 3 + \mathbb{E}\left[\mathrm{Geo}(4/n - 4/n^2)\right] \sim n/4$. If $c < 1/4$ then

$$\mathbb{P}\left[3 + \mathrm{Geo}(4/n - 4/n^2) \geq cn \log n\right] \geq (1 - 4/n + 4/n^2)^{cn \log n}\,.$$

Since $e^{-y} \geq 1 - y \geq e^{-y-y^2}$ for every $y \in [0, 1/4]$,

$$(1 - 4/n + 4/n^2)^{cn \log n} = \exp((-4/n + O(1/n^2))cn \log n) = (1/n)^{4c} e^{o(1)} \geq 1/n\,.$$

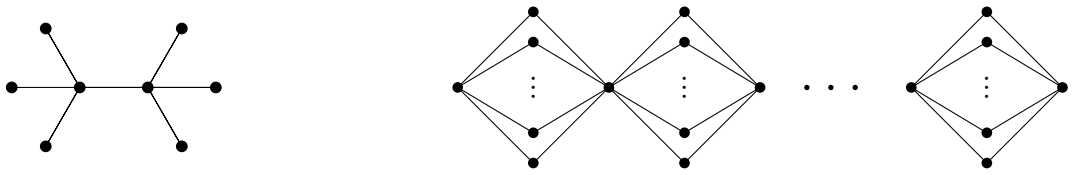

Figure 1: Left: the double star graph DS_8, which has a large guaranteed spread time in both variants. Right: a string of diamonds, on which the asynchronous push&pull protocol is much quicker than its synchronous variant.

While, if $c > 1/4$, then

$$\mathbb{P}\left[3 + \mathrm{Geo}\left(\frac{4}{n} - \frac{4}{n^2}\right) \geq cn \log n\right] = \left(1 - \frac{4}{n} + \frac{4}{n^2}\right)^{cn \log n - 3} \leq \exp(-4c \log n + o(1)) < 1/n\,,$$

whence $\mathrm{gst}_{\mathsf{s}}(DS_n) \sim (n \log n)/4$.

2.5 The string of diamonds

Let m and $k \geq 2$ be positive integers, and let $\mathcal{S}_{m,k}$ be the 'string of diamonds' graph given in Figure 1(Right), where there are m diamonds, each consisting of k edge-disjoint paths of length 2 with the same end vetices. A vertex with degree greater than two is called a *hub*. There are $m+1$ hubs and km non-hubs, giving a total of $n = km + m + 1$ vertices. It turns out that in this graph the asynchronous push&pull protocol is much quicker than its synchronous variant.

Let us analyze the average spread times in the two protocols, starting with the asynchronous case. Proposition 13 gives that for each edge e,

$$T(e) \stackrel{d}{=} \mathrm{Exp}(1/2 + 1/k) \stackrel{s}{\leq} \mathrm{Exp}(1/2)$$

and that $\{T(e)\}_{e \in E}$ are independent. Between any two consecutive hubs there are k disjoint paths of length 2, so the communication time between them is stochastically dominated by $Z := \min\{Z_1, \ldots, Z_k\}$, where the Z_i are independent random variables equal in distribution to the sum of two independent $\mathrm{Exp}(1/2)$ random variables.

Lemma 15. *We have* $\mathbb{E}[Z] = O(1/\sqrt{k})$.

Proof. For any $t \geq 0$ we have

$$\mathbb{P}[Z > t] = \prod_i \mathbb{P}[Z_i > t] = \mathbb{P}[Z_1 > t]^k \leq \left(1 - \mathbb{P}[\mathrm{Exp}(1/2) \leq t/2]^2\right)^k = \left(2e^{-t/4} - e^{-t/2}\right)^k\,.$$

Thus, using the inequality $2e^{-t/4} - e^{-t/2} \leq e^{-t^2/64}$ valid for $t \in [0, 4]$, we find

$$\mathbb{E}[Z] = \int_0^\infty \mathbb{P}[Z > t]\,\mathrm{d}t \leq \int_0^4 e^{-kt^2/64}\mathrm{d}t + \int_4^\infty (2e^{-t/4})^k \mathrm{d}t \leq 8\sqrt{\pi/k} + \frac{2^{k+2}}{ke^k} = O(1/\sqrt{k})\,. \quad \square$$

By Lemma 15, the expected time for all the hubs to learn the rumour is $O(mk^{-1/2})$. Once all the hubs learn the rumour, a degree 2 vertex pulls the rumour in Exp(1) time and the expected value of the maximum of at most km independent Exp(1) variables is $O(\log km)$. So by linearity of expectation, $\mathrm{ast}_{\mathsf{a}}(G) = O(\log n + mk^{-1/2})$.

In the synchronous case, for any G we have $\mathrm{ast}_{\mathsf{s}}(G) \geq \mathrm{diam}(G)$. For this graph, we get $\mathrm{ast}_{\mathsf{s}}(G) \geq 2m$. Choosing $k = \Theta\left((n/\log n)^{2/3}\right)$ and $m = \Theta\left(n^{1/3}(\log n)^{2/3}\right)$ gives

$$\mathrm{ast}_{\mathsf{a}}(G) = O(\log n) \text{ and } \mathrm{ast}_{\mathsf{s}}(G) = \Omega(n^{1/3}(\log n)^{2/3}) .$$

This graph has $\mathrm{ast}_{\mathsf{s}}(G)/\mathrm{ast}_{\mathsf{a}}(G) = \Omega\left((n/\log n)^{1/3}\right)$ and is the example promised by Corollary 7.

3 Extremal spread times for push&pull protocols

In this section we prove Theorems 1 and 2.

3.1 Proof of (1) and its tightness

For a given $t \geq 0$, consider the protocol which is the same as push&pull except that, if the rumour has not spread to all vertices by time t, then the new process reinitializes. Coupling the new process with push&pull, we obtain for any $k \in \{0, 1, 2, \dots\}$ that

$$\mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(G, v) > kt\right] \leq \mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(G, v) > t\right]^k . \tag{15}$$

and

$$\mathbb{P}\left[\mathrm{ST}_{\mathsf{s}}(G, v) > kt\right] \leq \mathbb{P}\left[\mathrm{ST}_{\mathsf{s}}(G, v) > t\right]^k . \tag{16}$$

Combining (15) with

$$\mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(G, v) > e\mathbb{E}\left[\mathrm{ST}_{\mathsf{a}}(G, v)\right]\right] < 1/e,$$

which comes directly from Markov's inequality, we obtain

$$\mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(G, v) > e \log n\, \mathbb{E}\left[\mathrm{ST}_{\mathsf{a}}(G, v)\right]\right] < 1/n.$$

Since $\mathbb{E}\left[\mathrm{ST}_{\mathsf{a}}(G, v)\right] \leq \mathrm{ast}_{\mathsf{a}}(G)$ for all v, this gives the right-hand inequality in (1) directly from the definition of $\mathrm{gst}_{\mathsf{a}}$. This inequality is tight up to the constant factor, as the double star has $\mathrm{ast}_{\mathsf{a}}(DS_n) = \Theta(n)$ and $\mathrm{gst}_{\mathsf{a}}(DS_n) = \Theta(n \log n)$ (see Section 2.4).

To prove the left-hand inequality, let $\tau = \mathrm{gst}_{\mathsf{a}}(G)$ and let v be a vertex such that $\mathbb{E}\left[\mathrm{ST}_{\mathsf{a}}(G, v)\right] = \mathrm{ast}_{\mathsf{a}}(G)$. Then

$$\mathrm{ast}_{\mathsf{a}}(G) = \int_0^\infty \mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(G, v) > t\right] dt = \sum_{i \in \mathbb{N}_0} \int_{i\tau}^{(i+1)\tau} \mathbb{P}\left[\mathrm{ST}_{\mathsf{a}}(G, v) > t\right] dt \leq \sum_{i \in \mathbb{N}_0} \frac{\tau}{n^i}$$

by (15) with $t = \tau$. Hence $\mathrm{ast}_{\mathsf{a}}(G) \leq \tau/(1 - 1/n)$. This inequality is tight up to a constant factor, as the star has $\mathrm{ast}_{\mathsf{a}}(G_n^*) = \Theta(\mathrm{gst}_{\mathsf{a}}(G_n^*)) = \Theta(\log n)$ (see Section 2.2).

3.2 Proof of the right-hand bound in (2) and its tightness

We will actually prove this using pull operations only. Indeed we will show $\mathrm{ast}_{\mathsf{a}}^{\mathsf{pull}}(G) < 4n$, where the superscript pull means the 'pull only' protocol. Since the path has $\mathrm{ast}_{\mathsf{a}}^{\mathsf{pull}}(P_n) \geq \mathrm{ast}_{\mathsf{a}}(P_n) = \Theta(n)$ (see Section 2.3), this bound would be tight up to the constant factor.

The proof is by induction: we prove that when there are precisely m uninformed vertices, just b of which have informed neighbours (we call these b vertices the *boundary vertices*), the expected remaining time for the rumour to reach all vertices is at most $4m-2b$. The inductive step is proved as follows. Let I denote the set of informed vertices, B the set of boundary vertices, and R the set of the remaining vertices. Let $|B| = b$ and $|B|+|R| = m$. Let $d(v)$ denote the degree of v in G and, for a set S of vertices, let $d_S(v)$ count the number of neighbours of v in S. We consider two cases.

Firstly, suppose that there exists a boundary vertex v with $d_R(v) \geq d_B(v)$. We can for the next step ignore all calls from vertices other than v, so the process is forced to wait until v is informed before any other vertices. This clearly gives an upper bound on the spread time. The expected time taken for v to pull the rumour from vertices in I is

$$\frac{d(v)}{d_I(v)} = \frac{d_I(v)+d_R(v)+d_B(v)}{d_I(v)} \leq 1 + \frac{2d_R(v)}{d_I(v)} \leq 1 + 2d_R(v).$$

Once v is informed, the number of uninformed vertices decreases by 1, and the number of boundary vertices increases by $d_R(v) - 1$. The inductive hypothesis concludes this case since

$$1 + 2d_R(v) + 4(m-1) - 2(b + d_R(v) - 1) < 4m - 2b.$$

Otherwise, if there is no such v, then any boundary vertex v has a 'pulling rate' of

$$\frac{d_I(v)}{d_I(v)+d_R(v)+d_B(v)} \geq \frac{1}{1+d_R(v)+d_B(v)} \geq \frac{1}{2d_B(v)} \geq \frac{1}{2b}.$$

Since there are b boundary vertices, together they have a pulling rate of at least $1/2$, so the expected time until a boundary vertex is informed is at most 2. Once this happens, m decreases by 1 and b either does not decrease or decreases by 1, and the inductive hypothesis concludes the proof.

3.3 Proof of the left-hand bound in (2) and its tightness

In this section we show for any vertex v_0 of a graph G we have $\mathbb{E}\left[\mathrm{ST}_{\mathsf{a}}(G, v_0)\right] > (\log n)/5$. This is tight (up to the constant) as the star has $\mathrm{ast}_{\mathsf{a}}(G_n^*) = O(\log n)$ (see Section 2.2). We give an argument for an equivalent protocol, defined below.

Definition (Two-clock-per-edge protocol)**.** On every edge place two exponential clocks, one near each end vertex. All clocks are independent. On an edge joining vertices u and v, the clocks both have rate $\deg(u)^{-1} + \deg(v)^{-1}$. Note that this is the rate of calls along that edge, combined, from u and v (see Proposition 13). At any time that the clock near u on an edge uv rings, and v knows the rumour but u does not, the rumour is passed to u.

Proposition 16. *The two-clock-per-edge protocol is equivalent to the asynchronous push&pull protocol.*

Proof. Consider an arbitrary moment during the execution of the two-clock-per-edge protocol. Let I denote the set of informed vertices. For any edge uv with $u \in I$ and $v \notin I$, the rate of calls along uv is $\deg(u)^{-1} + \deg(v)^{-1}$. Moreover, the edges act independently. So, the behaviour of the protocol at this moment is exactly the same as that of the asynchronous push&pull protocol. Hence, the two protocols are equivalent. □

In view of Proposition 16, we may work with the two-clock-per-edge protocol instead. Let X_v be the time taken for the first clock located near v to ring. Then X_v is distributed as $\mathrm{Exp}(f(v))$ where $f(v) = 1 + \sum \deg(u)^{-1}$, the sum being over all neighbours u of v. Hence, $\sum f(v) = 2n$.

On the other hand, for a vertex $v \neq v_0$ to learn the rumour, at least one of clocks located near v must ring. Thus

$$\max\{X_v : v \in V(G)\backslash\{v_0\}\} \overset{s}{\leq} \mathrm{ST}_{\mathsf{a}}(G, v_0)\,.$$

Let $X = \max\{X_v : v \in V(G)\backslash\{v_0\}\}$. Hence to prove $\mathbb{E}\,[\mathrm{ST}_{\mathsf{a}}(G, v_0)] > (\log n)/5$ it suffices to show $\mathbb{E}\,[X] > (\log n)/5$.

Let $\tau = \log(n-1)/3$ and $A = V(G) \setminus \{v_0\}$, Then we have

$$\begin{aligned}\mathbb{P}\,[X < \tau] &= \prod_{v\in A}(X_v < \tau) = \prod_{v\in A}(1 - e^{-\tau f(v)}) \leq \exp\left(-\sum_{v\in A} e^{-\tau f(v)}\right)\\ &\leq \exp\left(-(n-1)e^{-\tau\sum_v f(v)/(n-1)}\right) \leq \exp\left(-(n-1)e^{-3\tau}\right) = e^{-1}\,.\end{aligned}$$

Here the first inequality follows from $1 - x \leq e^{-x}$, the second from the arithmetic-geometric mean inequality, and the last one from $2n = \sum_v f(v) \leq 3(n-1)$ which holds for $n \geq 3$. Consequently,

$$\mathbb{E}\,[X] \geq \mathbb{P}\,[X \geq \tau]\,\tau \geq (1 - e^{-1})\log(n-1)/3 > \frac{1}{5}\log n$$

for all large enough n.

3.4 Proof of (3) and its tightness

The bounds in (3) follow immediately from (1) and (2). The left-hand bound is tight as the star has $\mathrm{gst}_{\mathsf{a}}(G_n^*) = \Theta(\log n)$ (see Section 2.2), and the right-hand bound is tight as the double star has $\mathrm{gst}_{\mathsf{a}}(DS_n) = \Theta(n \log n)$ (see Section 2.4).

3.5 Proof of (5) and its tightness

In this section we will prove $\mathrm{ast}_{\mathsf{s}}(G) < 4.6n$, which is tight up to the constant factor, as the path has diameter $n-1$ and hence $\mathrm{ast}_{\mathsf{s}}(P_n) \geq n - 1$.

The proof is similar to the one for the right-hand bound in (2) given in Section 3.2. Let $\alpha = \sqrt{e}/(\sqrt{e} - 1)$. We consider the 'pull only' protocol, and will prove inductively that when there are m uninformed vertices and b boundary vertices, the expected remaining time for the rumour to

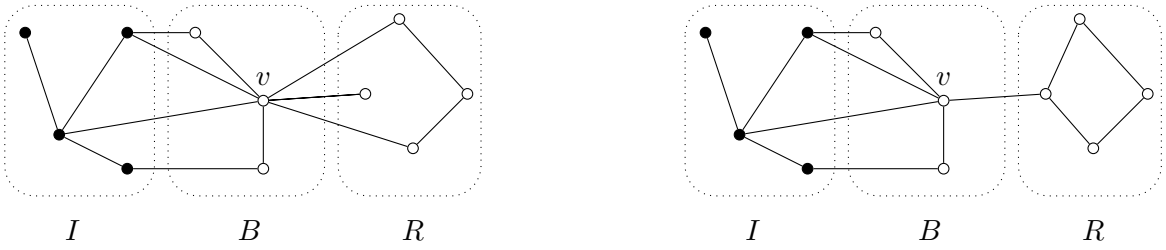

Figure 2: Left: first case in the proof of (5): there exists a boundary vertex v with $3 = d_R(v) \geq d_B(v) = 2$. Right: second case in the proof of (5): for all boundary vertices v we have $d_R(v) < d_B(v)$ (informed vertices are black, uninformed vertices are white).

reach all vertices is at most $(2+\alpha)m - 2b$, and it follows that $\mathrm{ast}_{\mathfrak{s}}(G) < 4.6n$. The inductive step is proved as follows. Let I denote the set of informed vertices, B the set of boundary vertices, and R the set of the remaining vertices. Let $|B| = b$ and $|B| + |R| = m$. Let $d(v)$ denote the degree of v in G and, for a set S of vertices, let $d_S(v)$ denote the number of neighbours of v in S. Consider two cases.

Firstly, suppose that there is a vertex $v \in B$ such that $d_R(v) \geq d_B(v)$ (see Figure 2(Left)). In this case, for the next step, we ignore all calls from vertices other than v and wait until v is informed before any other uninformed vertex. This gives an upper bound on the spread time. The expected time taken for v to pull the rumour from vertices in I is

$$1 + \mathbb{E}\left[\mathrm{Geo}\left(\frac{d_I(v)}{d(v)}\right)\right] = \frac{d(v)}{d_I(v)} = \frac{d_I(v) + d_R(v) + d_B(v)}{d_I(v)} \leq 2d_R(v) + 1.$$

Once v is informed, the number of uninformed vertices decreases by 1 and the number of boundary vertices increases by $d_R(v) - 1$. By the inductive hypothesis the expected time for the spread of the rumour is at most

$$2d_R(v) + 1 + (\alpha+2)(m-1) - 2(b + d_R(v) - 1) < (\alpha+2)m - 2b.$$

Next consider the case that $d_R(v) < d_B(v)$ for all $v \in B$ (see Figure 2(Right)). For each boundary vertex v we have

$$\frac{d_I(v)}{d(v)} = \frac{d_I(v)}{d_I(v) + d_R(v) + d_B(v)} \geq \frac{1}{1 + d_R(v) + d_B(v)} \geq \frac{1}{2d_B(v)} \geq \frac{1}{2b}.$$

Let X denote the time taken until the next vertex is informed. Then we have

$$X = 1 + \min\{X_1, X_2, \ldots, X_b\}\,,$$

where the X_i's are geometric random variables with parameters at least $1/2b$, and correspond to the waiting times of the boundary vertices, and they are independent since we are considering pull operations only. Thus we have

$$\mathbb{E}[X-1] = \sum_{t\in\mathbb{N}} \mathbb{P}[X - 1 \geq t] = \sum_{t\in\mathbb{N}} \prod_{i\in[b]} \mathbb{P}[X_i \geq t] \leq \sum_{t\in\mathbb{N}} \left(1 - \frac{1}{2b}\right)^{tb} \leq \sum_{t\in\mathbb{N}} e^{-t/2} = \alpha - 1\,,$$

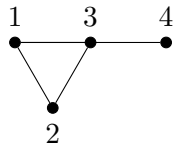

calling list for vertex 1:	2	2	3	2	3	$\cdots$
calling list for vertex 2:	1	3	3	1	3	$\cdots$
calling list for vertex 3:	2	2	1	4	1	$\cdots$
calling list for vertex 4:	3	3	3	3	3	$\cdots$

Figure 3: a particular outcome of the collection of calling lists for vertices

so a boundary vertex learns the rumour after at most α units of time on average, at which time the number of boundary vertices either does not decrease or decreases by 1. By inductive hypothesis again, the average spread time is at most

$$\alpha + (\alpha + 2)(m - 1) - 2(b - 1) = (\alpha + 2)m - 2b,$$

which completes the proof.

3.6 Proof of (4) and (6) and their tightness

The proof of (4) and its tightness are exactly the same as that for (1). The bound $\mathrm{gst}_{\mathsf{s}}(G) < 4.6en \log n$ is a direct consequence of bounds (5) and (4). This bound is tight (up to the constant factor) as the double star has guaranteed spread time $\Theta(n \log n)$ (see Section 2.4).

4 Comparison of the two protocols

We first prove Corollary 7 assuming Theorems 4 and 6, and in the following subsections we prove these theorems. The left-hand bound in Corollary 7 follows from Theorem 4; it is tight, up to the constant factor, as the star has $\mathrm{gst}_{\mathsf{a}}(G_n^*) = \Theta(\log n)$ and $\mathrm{gst}_{\mathsf{s}}(G_n^*) = 2$ (see Section 2.2). The right-hand bound in Corollary 7 follows from Theorem 6 by choosing $\alpha = 1/3$. A graph G was given in Section 2.5 having $\mathrm{ast}_{\mathsf{s}}(G)/\mathrm{ast}_{\mathsf{a}}(G) = \Omega\left((n/\log n)^{1/3}\right)$. Using (1) and (4), we get $\mathrm{gst}_{\mathsf{s}}(G)/\mathrm{gst}_{\mathsf{a}}(G) = \Omega\left(n^{1/3}(\log n)^{-4/3}\right)$ for this G.

4.1 The lower bound

In this section we prove Theorem 4. Let G be an n-vertex graph and let s denote the vertex starting the rumour. We give a coupling between the two versions. Consider a 'collection of calling lists for vertices': for every vertex u, we have an infinite list of vertices, each entry of which is a uniformly random neighbour of u, chosen independently from other entries, see Figure 3 for an example.

The coupling is built by using the same collection of calling lists for the two versions of the push&pull protocol. Note that $\mathrm{ST}_{\mathsf{s}}(G, s)$ is determined by this collection, but to determine $\mathrm{ST}_{\mathsf{a}}(G, s)$ we also need to know the Poisson processes associated with the vertices.

We first prove that $\mathrm{ast}_{\mathsf{a}}(G) \le H_n \times \mathrm{ast}_{\mathsf{s}}(G)$ (recall that H_n denotes the nth harmonic number). Consider the asynchronous protocol and let X_1 be the first time such that all clocks have rung

during the time interval $[0, X_1]$. Let X_2 be the first time such that all clocks have rung during the time interval $(X_1, X_1 + X_2]$, and define $X_3, X_4, \ldots$ similarly. Partition $[0, \infty)$ into subintervals $[0, X_1]$, $(X_1, X_1 + X_2]$, $(X_1 + X_2, X_1 + X_2 + X_3]$ etc. Consider a 'decelerated' variant $\mathcal{D}$ of the asynchronous push&pull protocol in which each vertex makes a call the first time its clock rings in each subinterval, but ignores later clock rings in that subinterval (if any). The spread time in $\mathcal{D}$ is stochastically larger than that in the asynchronous push&pull protocol, so without loss of generality we may and will work with $\mathcal{D}$. Coupling $\mathcal{D}$ and the synchronous protocol using the same calling lists and using induction gives

$$\mathrm{ST}_{\mathcal{D}}(G, s) \overset{s}{\leq} X_1 + X_2 + \cdots + X_{\mathrm{ST}_{\mathsf{s}}(G,s)} .$$

Since the X_i are i.i.d. and $\mathbb{E}[X_1] = H_n$, Wald's equation (see, e.g., [18, lemma 10.2.9]) gives

$$\mathbb{E}[\mathrm{ST}_{\mathsf{a}}(G, s)] \leq \mathbb{E}[\mathrm{ST}_{\mathcal{D}}(G, s)] \leq H_n \times \mathbb{E}[\mathrm{ST}_{\mathsf{s}}(G, s)] \leq H_n \times \mathrm{ast}_{\mathsf{s}}(G) ,$$

as required.

Next we prove that $\mathrm{gst}_{\mathsf{a}}(G) \leq 8\,\mathrm{gst}_{\mathsf{s}}(G) \log n$. Let B denote the event '$\mathrm{ST}_{\mathsf{s}}(G, s) \leq 2\,\mathrm{gst}_{\mathsf{s}}(G)$', which depends on the calling lists only. Inequality (16) gives $\mathbb{P}[B^c] \leq 1/n^2$. Partition the time interval $[0, 2\,\mathrm{gst}_{\mathsf{s}}(G) \times 4 \log n)$ into subintervals $[0, 4 \log n)$, $[4 \log n, 8 \log n)$, etc. Consider another 'decelerated' variant $\mathcal{D}'$ of the asynchronous push&pull protocol in which each vertex makes a call the first time its clock rings in each subinterval (if it does), but ignores later clock rings in that subinterval (if any). The spread time in $\mathcal{D}'$ is stochastically larger than that in the asynchronous push&pull protocol, so without loss of generality we may and will work with $\mathcal{D}'$. Let A denote the event 'during each of these $2\,\mathrm{gst}_{\mathsf{s}}(G)$ subintervals, all clocks ring at least once.' If A happens, then an inductive argument gives that for any $1 \leq k \leq 2\,\mathrm{gst}_{\mathsf{s}}(G)$, the set of informed vertices in the $\mathcal{D}'$ at time $4k \log n$ contains the set of informed vertices after k rounds of the synchronous version. Hence, if both A and B happen, then we would have

$$\mathrm{ST}_{\mathsf{a}}(G, s) \leq \mathrm{ST}_{\mathcal{D}'}(G, s) \leq (4 \log n)\,\mathrm{ST}_{\mathsf{s}}(G, s) \leq (8 \log n)\,\mathrm{gst}_{\mathsf{s}}(G) .$$

Hence to complete the proof, we need only show that $\mathbb{P}[A^c] \leq 1/n - 1/n^2$.

Let I denote a given subinterval. In the asynchronous version, the clock of any given vertex rings with probability at least $1 - n^{-4}$ during I. By the union bound, all clocks ring at least once during I, with probability at least $1 - n^{-3}$. The number of subintervals in the definition of A is $2\,\mathrm{gst}_{\mathsf{s}}(G)$, which is $O(n \log n)$ by (6). By the union bound again, $\mathbb{P}[A^c] = O\left(\log n / n^2\right)$, as required.

Theorem 4 is tight, up to the constant factors, as the star has $\mathrm{gst}_{\mathsf{a}}(G_n^*) = \Theta(\log n)$ and $\mathrm{gst}_{\mathsf{s}}(G_n^*) = 2$ (see Section 2.2).

4.2 The upper bound

In this section we prove Theorem 6. Fix $\alpha \in [0, 1)$. We want to prove

$$\mathrm{gst}_{\mathsf{s}}(G) \leq n^{1-\alpha} + 64 \mathrm{gst}_{\mathsf{a}}(G) n^{(1+\alpha)/2} . \tag{17}$$

We first sketch the proof. The main ingredients in the proof are a coupling between the two protocols, and sharp concentration bounds. Consider the asynchronous version. List the vertices in the order their clocks ring. The list ends once all the vertices are informed. Now consider the natural coupling between the two protocols, the synchronous actions follow the same ordering as in the list. We partition the list into blocks according to a certain rule in such a way that the blocks have the following property: the synchronous protocol in each round will inform a superset of the set of vertices informed by the asynchronous variant in any single block. For example, if we require that in each block each vertex communicates with the others at most once, then we would have this property. However, in order to get our bound, we need to use a more delicate rule for building the blocks. To conclude, we find an upper bound for the number of blocks, which coincides with the right-hand side of (17).

We now give the details. Let us fix an arbitrary starting vertex. Let $B_1, B_2, \dots$ be an i.i.d. sequence of vertices, where B_i is a uniformly random vertex of G. For each i, let W_i be a uniformly random neighbour of B_i, chosen independently of all other choices. Hence, $W_1, W_2, \dots$ is also an i.i.d. sequence of vertices (not necessarily having uniform distribution). We define a coupling between the two protocols by using the two sequences $(B_i)_{i\in\mathbb{N}}$ and $(W_i)_{i\in\mathbb{N}}$.

To define the coupled asynchronous scenario, we also need to know the ringing times of the clocks. Let $Z_1, Z_2, \dots$ be a sequence of i.i.d. exponentials with rate n (and mean $1/n$), and let this sequence be independent of $(B_i)_{i\in\mathbb{N}}$ and $(W_i)_{i\in\mathbb{N}}$. Then the coupled asynchronous scenario proceeds as follows: at time Z_1 the clock of vertex B_1 rings and it contacts W_1, then at time $Z_1 + Z_2$ the clock of B_2 rings and it contacts W_2, and so on.

We now define a third rumour spreading scenario, which corresponds to the so-called *sequential protocol* [28]. This protocol works as the asynchronous one except we put $Z_i = 1$ for all i. Hence, the scenario only depends on the sequences $(B_i)_{i\in\mathbb{N}}$ and $(W_i)_{i\in\mathbb{N}}$. Let N denote the first time that this protocol has informed all the vertices. Note that $N \geq n-1$ and N takes integral values. Observe that, in the asynchronous scenario, all vertices are informed right after N clocks have rung, and the spread time is $\sum_{i=1}^{N} Z_i$. The following lemma relates $\mathrm{gst}_{\mathsf{a}}(G)$ and N.

Lemma 17. *Define the event*

$$\mathcal{A} := \{N \leq 4n\,\mathrm{gst}_{\mathsf{a}}(G)\}\,.$$

Then we have $\mathbb{P}[\mathcal{A}] \geq 1 - O(1/n^2)$.

Proof. As the spread time of the asynchronous scenario is $\sum_{i=1}^{N} Z_i$, by definition of $\mathrm{gst}_{\mathsf{a}}$ we have $\mathbb{P}\left[\sum_{i=1}^{N} Z_i > \mathrm{gst}_{\mathsf{a}}(G)\right] \leq 1/n$. By (15) we have

$$\mathbb{P}\left[\sum_{i=1}^{N} Z_i > 2\,\mathrm{gst}_{\mathsf{a}}(G)\right] \leq 1/n^2\,.$$

So we need only show that

$$\mathbb{P}\left[N > 2\sum_{i=1}^{N} nZ_i\right] = O(1/n^2)\,.$$

Since Z_i's are i.i.d. exponentials with rate n, the random variables nZ_i are i.i.d. exponentials with rate 1, so for any fixed t, Cramér's Theorem gives

$$\mathbb{P}\left[\sum_{i=1}^{t} nZ_i < t/2\right] = e^{-ct}$$

for some positive constant c. Since $N \geq n-1$, we have

$$\mathbb{P}\left[N > 2\sum_{i=1}^{N} nZ_i\right] \leq \sum_{t=n-1}^{\infty} \mathbb{P}\left[N > 2\sum_{i=1}^{N} nZ_i \,\middle|\, N = t\right] = \sum_{t=n-1}^{\infty} e^{-ct} = \exp(-\Omega(n))\,,$$

as required. □

To define the coupled synchronous scenario, we need some definitions. For each vertex v, let $\pi(v)$ denote the probability that $W_j = v$. Recall that this probability does not depend on j. Call a vertex v *special* if $\pi(v) > n^{\alpha-1}$. Note that since $\sum \pi(v) = 1$, there are less than $n^{1-\alpha}$ special vertices.

We partition the list $B_1, W_1, B_2, W_2, \ldots$ into infinitely many finite *blocks* as follows. The first block is of the form

$$B_1, W_1, B_2, W_2, \ldots, B_j, W_j\,,$$

with j as large as possible, subject to the following conditions:

1. We have $B_i \notin \{B_1, W_1, \ldots, B_{i-1}, W_{i-1}\}$ for all $1 < i \leq j$.
2. If $W_i \in \{B_1, W_1, \ldots, B_{i-1}, W_{i-1}\}$ for some $1 < i \leq j$, then W_i is special.

Note that we choose the block to be as long as possible, hence we stop at W_j only if B_{j+1} already appears in $B_1, \ldots, W_j$, or W_{j+1} is non-special and it appears in $B_1, \ldots, W_j$, or both. If we have stopped at W_j, then a new block is started from B_{j+1}, and this process is iterated forever to define all the blocks. Note that each block has an even number of elements.

Let $S_1, S_2, \ldots$ denote the sizes of the blocks, and let N_b be the smallest number such that

$$S_1 + S_2 + \cdots + S_{N_b} \geq 2N\,.$$

The following lemma relates the spread time of the synchronous protocol and N_b.

Lemma 18. *The spread time of the synchronous push&pull protocol is stochastically smaller than* $N_b + n^{1-\alpha}$.

Proof. In this proof we only consider the finite list $B_1, W_1, \ldots, B_N, W_N$, which is partitioned into blocks as discussed before. We further split the blocks into smaller ones according to the following rule. Let v be a special vertex and assume that in the sequential scenario, it is informed exactly at time i. So, either $B_i = v$ or $W_i = v$. If the elements $B_i, W_i, B_{i+1}, W_{i+1}$ are contained in the same block, then we split this block at this point, putting everything up to B_i, W_i in one block and B_{i+1}, W_{i+1} and everything after in the other one. Since the number of splits equals the number

of special vertices, and there are less than $n^{1-\alpha}$ special vertices, the new total number of blocks is less than $N_b + n^{1-\alpha}$. We work with these refined blocks for the rest of the proof.

We couple with a modified version of the synchronous push&pull protocol, which we call the *lazy* scenario. We define the coupled lazy scenario inductively using the blocks. Assume that the kth block is

$$B_i, W_i, B_{i+1}, W_{i+1}, \dots, B_j, W_j \, .$$

Then in the kth round of the lazy scenario, vertex B_i contacts W_i, vertex B_{i+1} contacts W_{i+1} and so on, up until vertex B_j contacts W_j (all these communications happen at the same time). Moreover, a vertex that does not appear in this block, does not perform any action in the kth round. It is clear that stochastic upper bounds for the spread time of this lazy scenario carries over to the synchronous push&pull scenario.

To complete the proof we will show that the set of vertices informed by the lazy scenario after k rounds equals the set of vertices informed by the sequential scenario right after time $(S_1 + \cdots + S_k)/2$. (The factor of 2 appears here because a block with r communications has length $2r$.) The proof proceeds by induction. Assume that the kth block is

$$B_i, W_i, B_{i+1}, W_{i+1}, \dots, B_j, W_j \, .$$

If no repetition happens in this block at all, then it is clear that the lazy scenario in one round informs every vertex which the sequential one informs during times $i, i+1, \dots, j$. Notice the possible problem if a repetition happens: if during this block, x contacts y and tells her the rumour for the first time, and z also contacts y and asks her the rumour, then in the sequential scenario both y and z will learn the rumour by time j, whereas in the lazy scenario this is not the case because these operations happen at exactly the same time. However, if v is a repeated vertex in this block, then v is a special vertex, and moreover by the secondary splitting of the blocks, we know that it cannot be the case that v is informed in this block for the first time and appears again later in the block. Hence, no 'informing path' of length greater than one can appear in this block, and the proof is complete. □

Let $k = 64\,\mathrm{gst}_{\mathsf{a}}(G) n^{(1+\alpha)/2}$. The following lemma bounds N_b.

Lemma 19. *Define the event*

$$\mathcal{B} := \{S_1 + \cdots + S_k \geq 8\,\mathrm{gst}_{\mathsf{a}}(G) n\} \, .$$

Then we have $\mathbb{P}[\mathcal{B}] \geq 1 - O(1/n^2)$.

Before proving this lemma, let us see why it concludes the proof of Theorem 6. By Lemmas 17 and 19 and the union bound, with probability at least $1 - 1/n$ both events $\mathcal{A}$ and $\mathcal{B}$ happen. Assume this is the case. Then we have

$$S_1 + \cdots + S_k \geq 8\,\mathrm{gst}_{\mathsf{a}}(G) n \geq 2N \, ,$$

which means $N_b \le k$ by the definition of N_b. Together with Lemma 18, this implies that with probability at least $1 - 1/n$, the spread time of the synchronous push&pull protocol is at most $k + n^{1-\alpha}$, which gives (17).

Proof of Lemma 19. Let $\ell = n^{(1-\alpha)/2}/4$. We first show that

$$\mathbb{P}\left[S_1 > 2\ell\right] \ge 1/2\,. \tag{18}$$

Let j be arbitrary. We compute the conditional probability of $\{S_1 \ge 2j+2\}$ given that $\{S_1 \ge 2j\}$. On the event $\{S_1 \ge 2j\}$, the conditional probability that B_{j+1} is a repetition of a vertex already in the block $B_1, W_1, B_2, W_2, \ldots, B_j, W_j$ is $2j/n$. The probability that W_{j+1} is a repetition of a non-special vertex in the block is bounded above by $2jn^{\alpha-1}$, since there are at most $2j$ distinct vertices in the block so far, and W_{j+1} is a given non-special vertex with probability at most $n^{\alpha-1}$. So, we have

$$\mathbb{P}\left[S_1 \ge 2j+2 | S_1 \ge 2j\right] \ge 1 - 4jn^{\alpha-1} \ge \exp(-8jn^{\alpha-1})\,.$$

Consequently,

$$\mathbb{P}\left[S_1 > 2\ell\right] \ge \prod_{j=1}^{\ell} \exp(-8jn^{\alpha-1}) = \exp\left(-4n^{\alpha-1}\ell(\ell+1)\right) \ge 1/2$$

by the choice of ℓ, so (18) holds.

Observe that the block sizes $S_1, S_2, \ldots$ are i.i.d., and each of them is at least 2ℓ with probability at least $1/2$. So we have

$$\mathbb{P}\left[S_1 + \cdots + S_k \le k\ell/2\right] \le \mathbb{P}\left[\mathrm{Bin}(k, 1/2) \le k/4\right] \le \exp(-k/16) = O(1/n^2)\,,$$

where for the second inequality, we have used the Chernoff bound for binomials (see, e.g., [25, Theorem 2.3(c)]). □

References

[1] H. Amini, M. Draief, and M. Lelarge. Flooding in weighted sparse random graphs. *SIAM J. Discrete Math.*, 27(1):1–26, 2013.

[2] N. Berger, C. Borgs, J.T. Chayes, and A. Saberi. On the spread of viruses on the Internet. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 301–310, 2005.

[3] B. Bollobás and Y. Kohayakawa. On Richardson's model on the hypercube. In *Combinatorics, geometry and probability (Cambridge, 1993)*, pages 129–137. Cambridge Univ. Press, Cambridge, 1997.

[4] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.

[5] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. 6th Symp. Principles of Distributed Computing (PODC)*, pages 1–12, 1987.

[6] B. Doerr, M. Fouz, and T. Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. 43th Symp. Theory of Computing (STOC)*, pages 21–30, 2011.

[7] B. Doerr, M. Fouz, and T. Friedrich. Asynchronous rumor spreading in preferential attachment graphs. In *Proc. 13th Scandinavian Workshop Algorithm Theory (SWAT)*, pages 307–315, 2012.

[8] B. Doerr, M. Fouz, and T. Friedrich. Experimental analysis of rumor spreading in social networks. In *Design and analysis of algorithms*, volume 7659 of *Lecture Notes in Comput. Sci.*, pages 159–173. Springer, Heidelberg, 2012.

[9] R. Durrett. Stochastic growth models: recent results and open problems. In *Mathematical approaches to problems in resource management and epidemiology (Ithaca, NY, 1987)*, volume 81 of *Lecture Notes in Biomath.*, pages 308–312. Springer, Berlin, 1989.

[10] R. Elsässer and T. Sauerwald. On the runtime and robustness of randomized broadcasting. *Theoret. Comput. Sci.*, 410(36):3414–3427, 2009.

[11] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Struct. Algorithms*, 1(4):447–460, 1990.

[12] J. A. Fill and R. Pemantle. Percolation, first-passage percolation and covering times for Richardson's model on the n-cube. *Ann. Appl. Probab.*, 3(2):593–629, 1993.

[13] N. Fountoulakis and K. Panagiotou. Rumor spreading on random regular graphs and expanders. In *Proc. 14th Intl. Workshop on Randomization and Comput. (RANDOM)*, pages 560–573, 2010.

[14] N. Fountoulakis, K. Panagiotou, and T. Sauerwald. Ultra-fast rumor spreading in social networks. In *Proc. 23th Symp. Discrete Algorithms (SODA)*, pages 1642–1660, 2012.

[15] T. Friedrich, T. Sauerwald, and A. Stauffer. Diameter and broadcast time of random geometric graphs in arbitrary dimensions. *Algorithmica*, 67(1):65–88, 2013.

[16] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *28th International Symposium on Theoretical Aspects of Computer Science (STACS 2011)*, volume 9, pages 57–68, 2011.

[17] G. Giakkoupis. Tight bounds for rumor spreading with vertex expansion. In *Proc. 25th Symp. Discrete Algorithms (SODA)*, pages 801–815, 2014.

[18] G. R. Grimmett and D. R. Stirzaker. *Probability and random processes*. Oxford University Press, New York, third edition, 2001.

[19] M. Harchol-Balter, F. Thomson Leighton, and D. Lewin. Resource discovery in distributed networks. In *Proc. 18th Symp. Principles of Distributed Computing (PODC)*, pages 229–237, 1999.

[20] S. M. Hedetniemi, S. T. Hedetniemi, and A. L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.

[21] C. D. Howard. Models of first-passage percolation. In *Probability on discrete structures*, volume 110 of *Encyclopaedia Math. Sci.*, pages 125–173. Springer, Berlin, 2004.

[22] S. Janson. One, two and three times $\log n/n$ for paths in a complete graph with random weights. *Combin. Probab. Comput.*, 8(4):347–361, 1999.

[23] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized Rumor Spreading. In *Proc. 41st Symp. Foundations of Computer Science (FOCS)*, pages 565–574, 2000.

[24] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proc. 44th Symp. Foundations of Computer Science (FOCS)*, pages 482–491, 2003.

[25] C. McDiarmid. Concentration. In M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, and B. Reed, editors, *Probabilistic Methods for Algorithmic Discrete Mathematics*, pages 195–243. Springer-Verlag, 1998.

[26] A. Mehrabian and A. Pourmiri. Randomized rumor spreading in poorly connected small-world networks. *arXiv*, 1410.8175 [cs.SI], 2014. submitted (conference version in DISC 2014).

[27] K. Panagiotou and L. Speidel. Asynchronous rumor spreading on random graphs. In Leizhen Cai, Siu-Wing Cheng, and Tak-Wah Lam, editors, *Algorithms and Computation*, volume 8283 of *Lecture Notes in Computer Science*, pages 424–434. Springer Berlin Heidelberg, 2013.

[28] T. Sauerwald. On mixing and edge expansion properties in randomized broadcasting. *Algorithmica*, 56(1):51–88, 2010.